

A Comparison of Unified Storage for Enterprise AI

Tintri IntelliFlash versus NetApp and Dell EMC

by DCIG Lead Analyst / Storage, Ken Clipperton

PRODUCTS

Tintri IntelliFlash

URL ► <https://www.tintri.com/intelliflash>

Tintri by DDN
9351 Deering Ave
Chatsworth, CA 91311

NetApp FAS

URL ► <https://www.netapp.com/>

NetApp
650 Castro St #400
Mountain View, CA 94041
(800) 379-7873

Dell EMC Isilon

URL ► <https://www.delltechnologies.com>

Dell Technologies
1 Dell Way
Round Rock, TX 78664
Toll Free +1 (866) 438-3622

USE CASE

Storage systems for enterprise AI

The Growing Importance of Enterprise AI

The global demand for data-centric artificial intelligence (AI), analytics, and real-time insight is proliferating. Organizations increasingly seek to store multiple petabytes of structured and unstructured file data and employ new AI and machine learning (ML) processes to gain new value from that data. Storage to meet such requirements needs to be highly scalable, with flexibility to optimize multiple application datasets concurrently, along with intelligent management and analytics capabilities.

AI in the Enterprise

Many enterprises are now interested in running data intensive environments that, until recently, were mostly the domain of research institutions. Much of this development is driven by AI and ML applications that access data via file protocols. Thus, file-based workloads are at the heart of much AI-led innovation.

Datasets comprising both structured data—including call records and credit card transactions, as well as unstructured data—such as video, images and audio files, are also expanding as enterprises seek to gain new business value from consumer preference analysis and create products leveraging new levels of automation.

The growth of data and the need for real-time analytics in the enterprise is driving the demand for scalable, high-performance file storage optimized for AI workloads. Here are some key requirements:

- Scale, efficiency, and end-to-end optimization for an effective data strategy
- Powerful management tools and integrated enterprise features for data protection
- Flexible and versatile data movement between on-premises environments and public clouds

Tintri's IntelliFlash systems offer NVMe-accelerated performance and dense, cost-effective HDD-based hybrid storage capacity with concurrent multiprotocol support. These attributes extend the reach of IntelliFlash, meeting complex file service requirements and unlocking innovation in enterprise AI.

IntelliFlash systems have always supported both block and file protocols for virtualized environments. IntelliFlash now integrates robust and mature SMB and NFS capabilities, enabling IntelliFlash systems to deliver native file services for enterprise and HPC environments. These enhanced file server capabilities should

appeal to enterprises grappling with a mix of virtualized and physical hosts, alongside a rapid expansion of data, and new AI/ML/HPC workloads.

DDN Brings its HPC and At-Scale AI Infrastructure Expertise to the Enterprise

Tintri's parent company, DDN, is the leading storage provider for at-scale AI and HPC environments and is, in fact, the largest private storage company in the world. DDN has developed a deep understanding of data-intensive workloads and the demands they place on storage at any scale. While many enterprises are just now exceeding the petabyte-scale storage threshold, DDN has clients with exabyte-scale deployments.

Over the past several years, as enterprises grew more interested in the potential value of data-intensive AI, DDN grew more interested in applying its expertise in data at scale to new solutions for enterprise customers. DDN's acquisition of Tintri, Nexenta, and IntelliFlash product lines, now collectively under the Tintri brand, has opened up access to thousands of enterprises and their broader enterprise infrastructure ecosystems.

These acquisitions also brought into DDN some of the most advanced AI and ML technology in enterprise storage, bolstering its Intelligent Infrastructure capabilities. The broader DDN strategy is to expand these and other AI-related capabilities and autonomous operations across the whole range of DDN and Tintri products.

The AI-enabled Intelligent Infrastructure Opportunity

Most enterprise storage suppliers are embracing storage analytics and the proactive support that they make possible. The focus and maturity of vendor offerings vary. Some are focused on individual array health, relying more on fault data than predictive analytics. Others focus on application-level performance and availability. All offer real benefits to customers, including increased uptime and efficiency.

The end goal of predictive analytics for the more visionary storage providers goes beyond eliminating downtime. Their goal is to enable data center infrastructures to autonomously optimize themselves for application availability, performance, and total cost of ownership based on the customer's priorities. Tintri refers to these capabilities and attributes collectively as Intelligent Infrastructure.

This report compares Tintri IntelliFlash with NetApp and Dell EMC filers and their ability to support emerging enterprise AI workloads by evaluating the solutions based on the following four categories:

- Multiprotocol support to facilitate workload consolidation
- Consistent performance at scale with end-to-end data optimization
- Intelligent infrastructure analytics
- One consistent operating environment

Multiprotocol Support Enables Workload Consolidation

Concurrent, multiprotocol file and block support enables workload consolidation, which achieves several business benefits:

- Eliminates the need to acquire, manage, and maintain separate SAN and NAS storage environments
- Reduces acquisition costs, system management, and administrative overhead
- Minimizes data center costs for power, cooling, and rack space

Tintri IntelliFlash systems provide concurrent support for block (Fibre Channel and iSCSI) and file (NFS v3/v4, CIFS, and SMB3) protocols in a single system. This eliminates the need for separate SAN and NAS environments.

IntelliFlash systems provide the flexibility to optimize the key attributes demanded by an application. For example, an active/active controller configuration applies all resources (CPUs, memory, storage, and networking) to deliver maximum performance. Alternatively, an active/passive configuration will prioritize consistent performance under failure conditions.

IntelliFlash supports both physical and virtualized hosts. Its integrated intelligence and broad support for host operating systems, hypervisors, and applications combine the highly adaptable hardware resources, enabling organizations to consolidate many workloads onto a single IntelliFlash system.

NetApp FAS also provides concurrent, multiprotocol file and block support. NetApp also supports FC-NVMe on its all-flash AFF platforms that have FC ports and limited S3 object storage capabilities.¹

Like IntelliFlash, NetApp supports both active/active and active/passive controller configurations. Unlike IntelliFlash, however, NetApp customers must ensure that the same software options are licensed and enabled on both nodes. Otherwise, critical functionality required by some applications may be lost when component failures or routine maintenance results in a controller takeover.

Dell EMC Isilon supports file and S3 object protocols but not block protocols. Consequently, the Dell EMC solution can consolidate file and object workloads, but not the block workloads common in the enterprise. Thus, Dell EMC customers must maintain a separate storage infrastructure for the block-oriented SAN workloads, which introduces additional cost and complexity when compared to IntelliFlash.

The IntelliFlash advantage. Tintri IntelliFlash provides protocol support for the full spectrum of midrange enterprise AI workloads. Combined with the other advantages outlined below, it offers a compelling workload consolidation platform.

Consistent Performance at Scale and End to End Data Optimization

Tintri IntelliFlash scales file services without performance degradation through several innovations including patented distributed metadata acceleration, intelligent caching algorithms, and unique methods of virtualizing underlying storage for enhanced performance and rapid recovery from drive failures.

IntelliFlash employs a fully distributed storage architecture that logically separates data into three classes which vary by resiliency, performance, and capacity. IntelliFlash virtualizes the underlying storage media such that RAID groups are based on logical constructs rather than traditional disk-based RAID groups. IntelliFlash hybrid storage systems implement this architecture using NVDIMMs and NVMe for metadata and intelligent caching and HDDs for data storage. (Figure 1.)

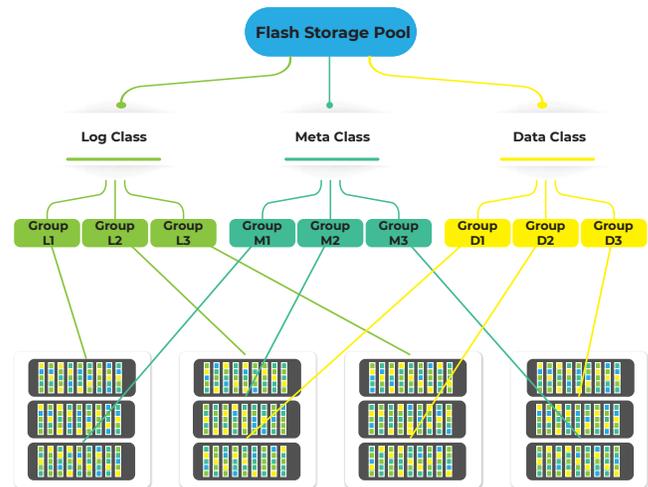


Figure 1

During normal operations, IntelliFlash distributes data slices across as many drives as required to achieve the necessary performance, resilience, and capacity for each logical volume. The system takes full advantage of the performance and capacity of all storage media and drives, intelligently distributing the workload to avoid creating hotspots.

IntelliFlash systems deduplicate and compress data as it is ingested, boosting the effective capacity of the already large caches. This accelerates all I/O by increasing the percentage of requests fulfilled from cache and reducing I/O activity on the back-end storage media.

Its intelligent caching algorithms continually self-adjust to deliver optimal performance from whatever mix of storage media is available. IntelliFlash leverages DRAM, NVDIMMs, NVMe-flash, and HDD media internally. It can also send snapshots to S3-based object storage. This enables IntelliFlash to cost-effectively deliver performance and capacity to meet a broad range of requirements. This is true of all IntelliFlash systems.

Tintri IntelliFlash delivers a highly-optimized, lockless data path that takes advantage of the parallelism that the NVMe protocol makes possible. In combination with the active-active, fully distributed IntelliFlash architecture, this NVMe-based parallelism unlocks the performance potential of all storage media to deliver scalable multi-petabyte file services at high performance.

1. <https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.dot-cm-sanconf%2FGUID-A39311D7-D917-4462-862C-4AC844EB4CA9.html>. Accessed 2020-12-07.

IntelliFlash delivers maximum performance from available resources during normal operations and rapidly recovers to normal operations after a drive failure.

Like IntelliFlash, **NetApp** FAS arrays employ NVDIMMs to coalesce writes. However, NetApp employs traditional RAID sets with hot spare drives. During normal operations, these spare drives do not contribute to array performance. In the event of a drive failure, the rebuild process can be lengthy because it is limited to that single hot spare drive's performance.

Dell EMC Isilon provides file services via a single file system and one namespace per scale-out cluster. Like IntelliFlash and NetApp FAS, Dell EMC Isilon nodes also use NVDIMMs to coalesce writes safely. Isilon aggregates the memory cache on each node in a cluster into one globally accessible pool of memory accessed over an internal 40Gb Ethernet cluster interconnect. Metadata is mirrored across multiple nodes, and data is striped across cluster nodes at the file level for redundancy.

The IntelliFlash advantage. The IntelliFlash storage architecture takes full advantage of all storage media, enabling high performance even as capacity scales up. Its efficient active/active dual-controller architecture leverages all system resources to deliver maximum performance in a footprint as small as two rack units. Its storage virtualization technology enables a rapid return to a fully protected state when media failures do occur. Neither the NetApp nor Dell EMC solutions offer this combination of advantages.

Intelligent Infrastructure Analytics

Cloud-based predictive analytics increase uptime by gathering and analyzing telemetry and system data such as capacity usage, configurations, system health, and performance. The results of these analyses are visible to both customers and support personnel.

Proactive support kicks in when system health alerts or the predictive analytics identify an issue. Administrators are notified. Cases are opened. Parts are shipped. Problems are resolved, in many cases without the downtime that would otherwise have occurred

Tintri IntelliFlash

Analytics for IntelliFlash delivers the deep intelligence and intuitive tools to maximize uptime and performance and keep systems operating at peak efficiency. It transforms storage administration by leveraging AI-driven analytics to monitor and optimize all IntelliFlash systems.

Analytics for IntelliFlash aggregates the collective user experience from thousands of customer installations worldwide to predict future requirements, provide proactive alerts, and detect problems before they develop into component and system failures. Intelligent analytics collects millions of data points from every IntelliFlash system deployed (where customers have opted to participate). Servers in the cloud process and analyze data to detect issues and identify patterns that can predict trends. Since this is all done in the cloud, there's no need to install any agents or stand up separate infrastructure.

While IntelliFlash analytics is an example of Tintri "Intelligent Infrastructure," it should be noted that the term refers to a spectrum of capabilities that span the Tintri portfolio. Having acquired several advanced storage analytics technologies, DDN has embarked on a purposeful cross-pollination of these intelligent infrastructure capabilities across all its storage systems. To facilitate this cross-pollination, DDN

has unified engineering for all products in a single development organization, taking advantage of each team's strengths and experiences.

NetApp has focused its cloud-based ActiveIQ analytics on system health, availability, and security. It makes recommendations based on those analytics. NetApp provides API services for integration into the change management frameworks embraced by many enterprises and Ansible playbooks for automation.

Dell EMC Isilon is on par with NetApp with its cloud-based CloudIQ. Currently, Dell Isilon administrators can—and may need to—manually optimize data and metadata layout per file or directory based on workload requirements.

The IntelliFlash advantage. Tintri IntelliFlash's infrastructure analytics compares favorably with NetApp and Dell EMC. Tintri's expertise in AI/ML is delivering rapid advances in infrastructure analytics across all DDN storage systems, enabling IT to focus on adding value to the business rather than on infrastructure administration.

One Consistent Operating Environment

Tintri IntelliFlash. From its very beginnings the IntelliFlash team has focused on revolutionizing the total cost of ownership for enterprise storage. To do so, they introduced innovations in both technology and business practices. For example, IntelliFlash was one of the first products to make inline deduplication and compression work for active workloads across its hybrid and all-flash systems. These innovations – available with a single OS that spans all IntelliFlash systems - reduce the cost of acquiring and owning IntelliFlash storage.

As a related point, IntelliFlash is also among the first storage platforms to offer customer-friendly all-inclusive software licensing, which simplifies the initial purchasing process by reducing the number of decision points. All-inclusive licensing also creates a consistent operating environment. With IntelliFlash, there is no need to figure out whether a particular feature is worth the extra cost for a particular storage system. Every feature is available to every IntelliFlash system on day one, **and** when business requirements suddenly change on day two.

NetApp FAS and AFF both run ONTAP, and NetApp is taking steps to reduce complexity with updates to its management software and more generous software bundles. Nevertheless, it still bases licensing on feature bundles that vary depending on the array model, and requires a-la-carte licenses for some features. NetApp even ties some feature licenses to specific nodes by serial number.

Dell EMC Isilon and PowerScale both run OneFS. Nevertheless, licensing and supported features vary by product. This variability adds complexity to the ownership process. For example, Isilon licensing is based on feature bundles plus a-la-carte licenses. Also, OneFS supports the use of self-encrypting drives on Isilon but not on PowerScale. PowerScale customers must learn about and plan around such differences.

The IntelliFlash advantage. All-inclusive feature licensing increases agility and reduces management overhead. It enables IT to be more responsive to the business by eliminating the costs and delays associated with justifying, approving, purchasing, installing, and managing feature licenses on an a-la-carte basis. Those a-la-carte processes can add weeks or even months of delay. With all-inclusive feature licensing, all the system's capabilities are immediately available to the business.

Why buy Tintri IntelliFlash over NetApp and Dell EMC for Enterprise AI Workloads?

- Nobody knows at-scale data management in production environments like DDN.
- IntelliFlash combines the performance, capacity, and protocol support to serve the full spectrum of midrange enterprise AI/ML/HPC workloads.
- IntelliFlash enables organizations to consolidate multiple SAN and NAS storage systems in one efficient, high-performance, high-density, low management overhead system that scales to multi-petabytes.
- Ongoing DDN/Tintri technology integration gives customers the best of both worlds. Tintri IntelliFlash combines proven DDN NVMe platforms designed to extract the most performance at scale with Tintri enterprise features and sophisticated storage analytics.
- One consistent operating environment with identical features across all IntelliFlash systems reduces management overhead and enhances agility.

KEY QUESTIONS TO ASK:

- Will your company need to store multiple petabytes of structured and unstructured data within the new storage system's lifetime?
- Is your company seeking to gain new value from its data through AI/ML/HPC applications?
- Is your company seeking to accelerate databases and other critical applications on-premises?
- Is your company more interested in using the cloud as a backup target than for running critical workloads?
- Does your company need a complete high-performance storage system it can deploy in just 2RU?

If the answer to any of these questions is "Yes," then Tintri IntelliFlash may be the best fit for your organization. ■

About DCIG

DCIG, the Data Center Intelligence Group, empowers the information technology industry with actionable analysis. DCIG provides informed third-party analysis of various cloud, data protection, and data storage technologies. Learn more at www.dcig.com.



DCIG, LLC // 7511 MADISON STREET // OMAHA NE 68127 // 844.324.4552

dcig.com

© 2021 DCIG, LLC. All rights reserved. Other trademarks appearing in this document are the property of their respective owners. This DCIG Competitive Intelligence Report Executive Edition is a product of DCIG, LLC. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. Product information was compiled from both publicly-available and vendor-provided resources. While DCIG has attempted to verify that product information is correct and complete, feature support can change and is subject to interpretation. All features represent the opinion of DCIG. No negative inferences should be drawn against any product or vendor not included in this report. DCIG cannot be held responsible for any errors that may appear.